

WOODHOUSE EXHIBIT F

EXHIBIT 9

Message

From: Guillaume Lample [REDACTED]@meta.com]
Sent: 1/4/2023 7:10:13 PM
To: Faisal Azhar [REDACTED]@meta.com]; Guillaume Lample [REDACTED]@meta.com]
Subject: Message summary [{"otherUserFbId":100076968606838,"threadFbId":null}]

Guillaume Lample (1/04/2023 05:53:36 PST):

>Hi Faisal

Guillaume Lample (1/04/2023 05:53:41 PST):

>do you know how we can ask for more storage ?

Guillaume Lample (1/04/2023 05:53:49 PST):

>we are really limited right now and this is totally blocking

Faisal Azhar (1/04/2023 05:55:40 PST):

>Hi Guillaume, How are you doing? I'll have to look into this, I assume this is for FAIR cluster? I have on my to do list to track getting back our previous storage on FAIR cluster.

Guillaume Lample (1/04/2023 05:55:52 PST):

>yes FAIR cluter

Guillaume Lample (1/04/2023 05:56:01 PST):

>but i was wondering whether we can ask for more

Guillaume Lample (1/04/2023 05:56:08 PST):

>I mean some projects have 350TB

Faisal Azhar (1/04/2023 05:56:15 PST):

>What's our current storage?

Faisal Azhar (1/04/2023 05:57:34 PST):

>I remember we need 25TB back

Guillaume Lample (1/04/2023 05:57:50 PST):

>75TB

Guillaume Lample (1/04/2023 05:57:56 PST):

>but lib-genesis is 70TB for instance

Guillaume Lample (1/04/2023 05:58:04 PST):

>and we want to pre-process a new version of common crawl

Guillaume Lample (1/04/2023 05:58:08 PST):

>which will be 50TB at least

Guillaume Lample (1/04/2023 05:58:27 PST):

>if some projects have 350TB I was thinking maybe we can ask 200TB overall ?

Faisal Azhar (1/04/2023 06:02:38 PST):

>I need some more information, e.g, is this storage capacity needed for Xlformer work or Theorem's math workstream? and how long will we need it for? - also is there any way to circumvent this need, e.g., using any other cluster storage etc. - Asking these as we will be questioned on the need.

Guillaume Lample (1/04/2023 06:03:25 PST):

>it will be for xlformers

Guillaume Lample (1/04/2023 06:03:30 PST):

>1 year i guess

Guillaume Lample (1/04/2023 06:03:39 PST):

>no there is no way to circumvent

Guillaume Lample (1/04/2023 06:04:00 PST):

>also need to point out that multigen is probably not doing any effort to reduce their usage

Guillaume Lample (1/04/2023 06:04:03 PST):

>350TB wtf

Faisal Azhar (1/04/2023 06:04:09 PST):

>I assume that;s max, not min, right?

Guillaume Lample (1/04/2023 06:05:06 PST):
>yeah max

Guillaume Lample (1/04/2023 06:05:08 PST):
>min 6 month

Faisal Azhar (1/04/2023 06:05:23 PST):
>btw we should sync up so that I can understand your plans and work out needs to action.

Guillaume Lample (1/04/2023 06:05:39 PST):
>of course

Faisal Azhar (1/04/2023 06:05:47 PST):
>I'll put a meeting invite, would you prefer today or later this week?

Guillaume Lample (1/04/2023 06:05:56 PST):
>today is fine

Guillaume Lample (1/04/2023 06:06:18 PST):
>few meetings this afternoon though

Guillaume Lample (1/04/2023 06:06:40 PST):
>was supposed to have one at 3pm but the person is not showing up :D

Guillaume Lample (1/04/2023 06:06:43 PST):
>so now if you want

Faisal Azhar (1/04/2023 06:07:53 PST):
>Oh ok, let me send you an updated invite

Faisal Azhar (1/04/2023 06:09:10 PST):
>Ok see if you can join with the zoom link:
<https://fb.zoom.us/j/96653162394?pwd=eHAYWHdGcytTVlltY3lkU05wUm40QT09>

Guillaume Lample (1/04/2023 06:19:52 PST):
><https://docs.google.com/spreadsheets/d/> [REDACTED]

Faisal Azhar (1/04/2023 06:21:05 PST):
><https://docs.google.com/document/d/> [REDACTED] edit

Faisal Azhar (1/04/2023 09:08:48 PST):
>QQ: can we use S3 for data storage for pre-processing instead of storage space on FAIR cluster?

Faisal Azhar (1/04/2023 10:03:03 PST):
>Hey Guillaume, So I found the process to ask for storage but there is no storage available so that will require prioritisation among project.

>
>I also got a few questions in return that I need your help with.
>1. Can we run LLM on FAIR cluster v100?
>2. Can we use RSC instead to store data and run training of LLM? - hundreds of petabytes of flasharray available on RSC

Guillaume Lample (1/04/2023 10:04:15 PST):
>1) we can do inference yes, training too slow

Guillaume Lample (1/04/2023 10:04:35 PST):
>2) training will be done on RSC no matter what, but data have to sit on H2 at some point during preprocessing

Faisal Azhar (1/04/2023 10:05:07 PST):
>so thats a No for training on FAIR

Faisal Azhar (1/04/2023 10:05:44 PST):
>so why can not we simply put data on RSC? Alternative, what requires us to keep data on H2?

Guillaume Lample (1/04/2023 10:05:54 PST):
>yes finetuning is possible but absolutely not pretraining

Guillaume Lample (1/04/2023 10:06:01 PST):
>there is no internet on RSC

Guillaume Lample (1/04/2023 10:06:07 PST):
>cant download

Faisal Azhar (1/04/2023 10:06:18 PST):
>but we can copy data onto rsc via rsc tool

Guillaume Lample (1/04/2023 10:06:32 PST):

>sure but we need to download it and process it

Guillaume Lample (1/04/2023 10:06:37 PST):
>download is torrenting

Guillaume Lample (1/04/2023 10:06:42 PST):
>cant be done on RSC

Guillaume Lample (1/04/2023 10:06:46 PST):
>and devserver has tiny storage

Faisal Azhar (1/04/2023 10:07:18 PST):
>Got it, how about using s3 to download and then move to RSC? or if someone has already downloaded common crawl then copy onto RSC

Guillaume Lample (1/04/2023 10:07:18 PST):
>cant we argue about the insane usage of 350B of multigen ?

Guillaume Lample (1/04/2023 10:07:23 PST):
>which has been decelerated

Faisal Azhar (1/04/2023 10:07:33 PST):
>hahaa I am not sure I am in a position to do so lol

Guillaume Lample (1/04/2023 10:07:40 PST):
>i mean we cant process on S3

Faisal Azhar (1/04/2023 10:08:25 PST):
>how about process in RSC after moving from s3, is that not possible?

Guillaume Lample (1/04/2023 10:11:21 PST):
>at some point all the raw data has to be on a single cluster which is accessible by code

Guillaume Lample (1/04/2023 10:11:24 PST):
>so it cant be s3

Guillaume Lample (1/04/2023 10:11:37 PST):
>because code cannot process stuff on s3, and it cant be RSC because RSC has no internet

Guillaume Lample (1/04/2023 10:11:52 PST):
>so no , there are no other solutions

Guillaume Lample (1/04/2023 10:12:10 PST):
>AWS would have been a solution we have a little bit of storage there but since we are going to lose it soon...

Guillaume Lample (1/04/2023 10:12:15 PST):
>do you know when we will lose it exactly ?

Guillaume Lample (1/04/2023 10:12:20 PST):
>if it's end of february, that's nice

Guillaume Lample (1/04/2023 10:12:31 PST):
>have time to process stuff and transfer the processed version to RSC

Faisal Azhar (1/04/2023 10:23:41 PST):
>Re AWS, I'll check and circle back to confirm.

Faisal Azhar (1/04/2023 11:05:15 PST):
>Sorry for my confusion, I am still trying to understand why we cannot download all raw data on S3, move it to rsc using rsc tool, and then process stuff on RSC. Am I missing something here?

Guillaume Lample (1/04/2023 11:05:52 PST):
>We can't transfer from s3 to rsc

Guillaume Lample (1/04/2023 11:06:03 PST):
>We need to go through de server

Guillaume Lample (1/04/2023 11:06:07 PST):
>Devserver

Guillaume Lample (1/04/2023 11:06:13 PST):
>which has almost no storage

Faisal Azhar (1/04/2023 11:06:19 PST):
>aah I see

Guillaume Lample (1/04/2023 11:06:24 PST):

>Plus, transfer on rsc is insanely slow

Guillaume Lample (1/04/2023 11:06:38 PST):

>So we prefer to transfer only the cleaned processed dataset

Guillaume Lample (1/04/2023 11:06:47 PST):

>Also we have quite a limited storage on rsc too

Guillaume Lample (1/04/2023 11:06:56 PST):

>Unless we use airstore which is a nightmare

Guillaume Lample (1/04/2023 11:07:19 PST):

>We need to do something about multigen and these projects

Faisal Azhar (1/04/2023 11:08:17 PST):

>I am concerned that Patrick/DINO are going back on the agreement to return 25TB, what do you think?
Mixing the other project storage with our return complicates things and is an easy way to say NO to us.

Guillaume Lample (1/04/2023 11:09:42 PST):

>Yes Dino should stick to what they said

Guillaume Lample (1/04/2023 11:09:45 PST):

>That was the deal

Faisal Azhar (1/04/2023 11:10:13 PST):

>I agree